Image enhancement, image quality, and noise

Zia-ur Rahman[†], Daniel J. Jobson[‡], Glenn A. Woodell[‡], Glenn D. Hines[‡]

[†]College of William & Mary, Department of Applied Science, Williamsburg, VA 23187.
[‡]NASA Langley Research Center, Hampton, Virginia 23681.

ABSTRACT

The Multiscale Retinex With Color Restoration (MSRCR) is a non-linear image enhancement algorithm that provides simultaneous dynamic range compression, color constancy and rendition. The overall impact is to brighten up areas of poor contrast/lightness but not at the expense of saturating areas of good contrast/brightness. The downside is that with the poor signal-to-noise ratio that most image acquisition devices have in dark regions, noise can also be greatly enhanced thus affecting overall image quality. In this paper, we will discuss the impact of the MSRCR on the overall quality of an enhanced image as a function of the strength of shadows in an image, and as a function of the root-mean-square (RMS) signal-to-noise (SNR) ratio of the image.

1. INTRODUCTION

In the mid to late 1990s, Jobson, et al.[1,2] resurrected Edwin Land's retinex[3–5] and developed a general purpose image enhancement algorithm that they called the Multiscale Retinex with Color Restoration (MSRCR). Whereas the retinex was conceived by Land as a model of the lightness and color perception of human vision,[4] the MSRCR is neither defined nor used as a model for human vision color constancy. Rather, it is used as a digital image enhancement platform for synthesizing local contrast improvement, color constancy, and lightness/color rendition. The intent is to transform the visual characteristics of the recorded digital image so that the rendition of the transformed image approaches that of the direct observation of scenes. Special emphasis is placed on increasing the local contrast in dark zones of the recorded image so that it would match our perception of wide dynamic range scenes, e.g., scenes which contain objects that are partly in sunlight and partly in shadow.

Generally, images can be grouped into two broad categories: those that have sufficient signal-to-noise ratio (SNR), and those that do not. Images that fall in the first category are typically those with levels of brightness and contrast that are generally classified as good—we will discount out-of-focus or blurred images from the discussion. Conversely, images in the second category tend to be dark with insufficient brightness and contrast. Images in the second category also tend to be the images that benefit the most from image enhancement. The impact of the MSRCR on images as a function of their SNR is the primary focus of this paper. The MSRCR's local contrast enhancement typically results in enhancement of fine changes in spatial and spectral detail: i.e., it brings up low contrast features where they are obscured due to low intensity values. Since the MSRCR has no internal mechanism for differentiating noise from signal—a topic worthy of a PhD dissertation and more in its own right—the noise, as well as the signal, gets enhanced, especially in areas of the image with poor brightness, i.e, low intensity values. Thus, the overall impact of the MSRCR on images with areas of poor SNR is to enhance the noise features leading to an output image that can, at times, look extremely noisy. However, since this noise enhancement is often accompanied by greatly increased visibility of features in the dark region, the degree to which the MSRCR enhanced noise is objectionable can only be determined by the application. Regardless of the application, though, the impact of the MSRCR on images with poor SNR is to make the noise "visible," thus compromising image quality according to most traditional image analysis criteria. One way to reduce the impact of MSRCR on noise is to perform noise reduction. Noise reduction can be performed prior to the application of the MSRCR or afterwords. Since the MSRCR is a non-linear operator, the order of application impacts the final image quality. We will investigate both situations using a noise-reduction method[6] that attempts to compensate for noise without affecting the real edges in an image (Section 6), and present some preliminary results.

ZR: zrahman@as.wm.edu; DJJ: daniel.j.jobson@nasa.gov; GAW: glenn.a.woodell@nasa.gov; GDH: glenn.d.hines@nasa.gov

2. THE MULTI-SCALE RETINEX

The basic form of the Multiscale retinex (MSR) is given by[2,7,8]

$$R_i(x_1, x_2) = \sum_{k=1}^{K} W_k R_{k_i}, \quad i = 1, \ldots, N \tag{1}$$

$$R_{k_i} = \log\ I_i(x_1, x_2) - \log\left[F_k(x_1, x_2) * I_i(x_1, x_2)\right] \tag{2}$$

where R_k is the single-scale Retinex (SSR) defined in Equation 2, the sub-index i represents the i^{th} spectral band, N is the number of spectral bands—$N = 1$ for grayscale images and $N = 3$ for typical color images—I is the input image, R is the MSR output, F_k represents the k^{th} surround function, W_k is the weight associated with F_k, K is the number of surround functions, or scales, and '$*$' represents the convolution operator. The surround functions, F_k are given as:

$$F_k(x_1, x_2) = \kappa \exp[-(x_1^2 + x_2^2)/\sigma_k^2], \quad k = 1, \ldots, K, \tag{3}$$

$$\kappa = 1/\left(\sum_{x_1}\sum_{x_2} F(x_1, x_2)\right), \tag{4}$$

where σ_k are the scales that control the extent of the surround—smaller values of σ_k lead to narrower surrounds.

Each SSR is capable of enhancing some particular characteristic of the input image.[2,8] For instance, narrow surrounds highlight the fine features in the image, Figure 1(b), but almost all tonal rendition is lost. Wide surrounds, Figure 1(d) retain all the tonal information but do not enhance the small fine features. Scales in between the two extremes—medium scales—tend to enhance some small features and retain considerable tonal information, Figure 1(c), but typically lack overall tonal rendition. Hence, multiple surrounds are needed to achieve a graceful balance between dynamic range compression and tonal rendition. The number of scales used for the MSR is, of course, application dependent. We have found empirically, however, that a combination of three scales representing narrow, medium, and wide surrounds is sufficient to provide both dynamic range compression and tonal rendition. The MSR output in Figure 1(e) is formed with $\sigma_k = 15, 80, 250$, and $W_k = 1/3$ for $K = 3$. It enhances the spatial details in the original image, Figure 1(a), yet retains the nice tonal rendition from the output of a wide-surround SSR. The original image was 512×512 pixels.

3. MSR WITH COLOR RESTORATION

The general effect of the MSR on images that have significant areas of uniform or slowly changing intensities is a "graying out" of those areas. This desaturation of color can be severe depending upon the extent of the uniform area. In order to compensate for this desaturation, we developed a *color restoration* operator that seeks to produce good color rendition for images with any degree of graying. In addition this operator preserves a reasonable degree of color constancy. Since color constancy is known to be imperfect in human visual perception, some level of illuminant color dependency is acceptable provided it is much lower than the physical spectrophotometric variations. Ultimately this is a matter of image quality and color dependency is tolerable to the extent that the visual defect is not visually too strong.

In developing the color restoration operator, we began from the foundations of colorimetry[9] even though this is often considered to be in direct opposition to color constancy models. It is felt that colorimetry describes only the so-called "aperture mode" of color perception, i.e. restricted to the perception of color lights rather than color surfaces.[10] We made this choice simply because it serves as a foundation for creating a *relative* color space and in doing so, uses ratios that are less dependent on illuminant spectral distributions than the raw spectrophotometry. The color restoration factor, α, is computed based on the following transform:

$$\alpha_i(x_1, x_2) = f\left(I_i(x_1, x_2)/\sum_{n=1}^{N} I_n(x_1, x_2)\right), \quad i = 1, \ldots, N, \tag{5}$$

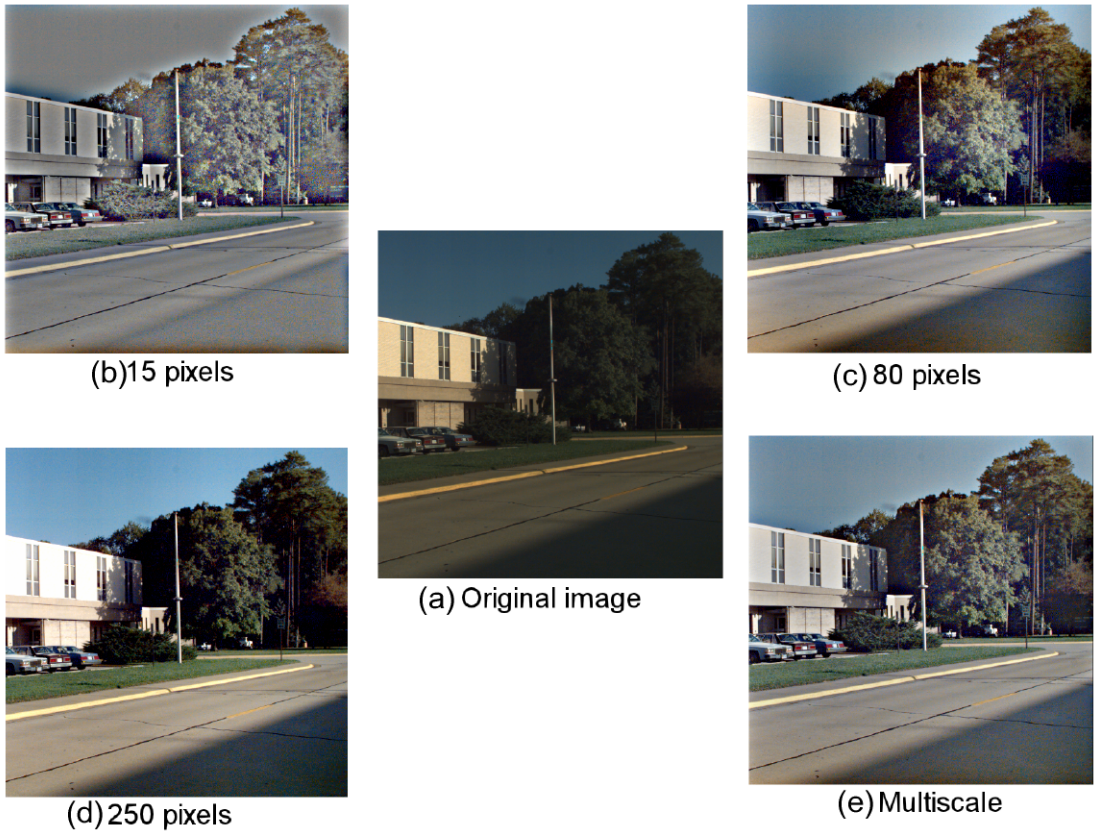

Figure 1. (a) The original input (b) Narrow surround (c) Medium surround (d) Wide surround (e) MSR output. The narrow-surround acts as a high-pass filter, capturing all the fine details in the image but at a severe loss of tonal information. The wide-surround captures all the fine tonal information but at the cost of dynamic range. The medium-surround captures both dynamic range and tonal information. The MSR is the average of the three renditions.

where $\alpha_i(x_1, x_2)$ is the color restoration coefficient in the i^{th} spectral band, N is the number of spectral bands, I_i is the i^{th} spectral band in the input image, and $f(\,)$ is some mapping function. In a purely empirical fashion we tested several renditions of $f(\,)$ on several retinexed images to gain a sense of their visual impact. The best results were obtained with a spectral analog of the SSR, i.e, where $f(\,) = \log(\,)$. This proved to restore color rendition, encompassing both saturated and less saturated colors. Adding this to equation 1, the Multiscale Retinex With Color Restoration (MSRCR) is given by:

$$R_i(x_1, x_2) = \alpha_i(x_1, x_2) \sum_{k=1}^{K} W_k \left(\log\ I_i(x_1, x_2) - \log\left[F_k(x_1, x_2) * I_i(x_1, x_2)\right]\right) \quad i = 1, \ldots, N. \tag{6}$$

The results of applying this transformation to the "monochrome" images are shown in Figure 2. The top row shows a set of images, each one of which violates the gray world assumption in some way. In addition to areas of uniform or slowly varying intensities, these violations also encompass scene features where there is an obvious color opponency/distortion. While the second and third columns of Figure 2 contain images that have large areas of slowly varying intensities, the first and the last columns have more to do with opponent color and slowly changing illuminations. The MSR output brings out the fine detail and produces images with better contrast. However, there is a significant desaturation of color. The MSRCR output shows, though, that the introduction of the color restoration function does indeed restore the color in the output images. There is an accompanying loss of contrast but, in most cases, the improved color more than compensates for the slight loss in contrast and sharpness.

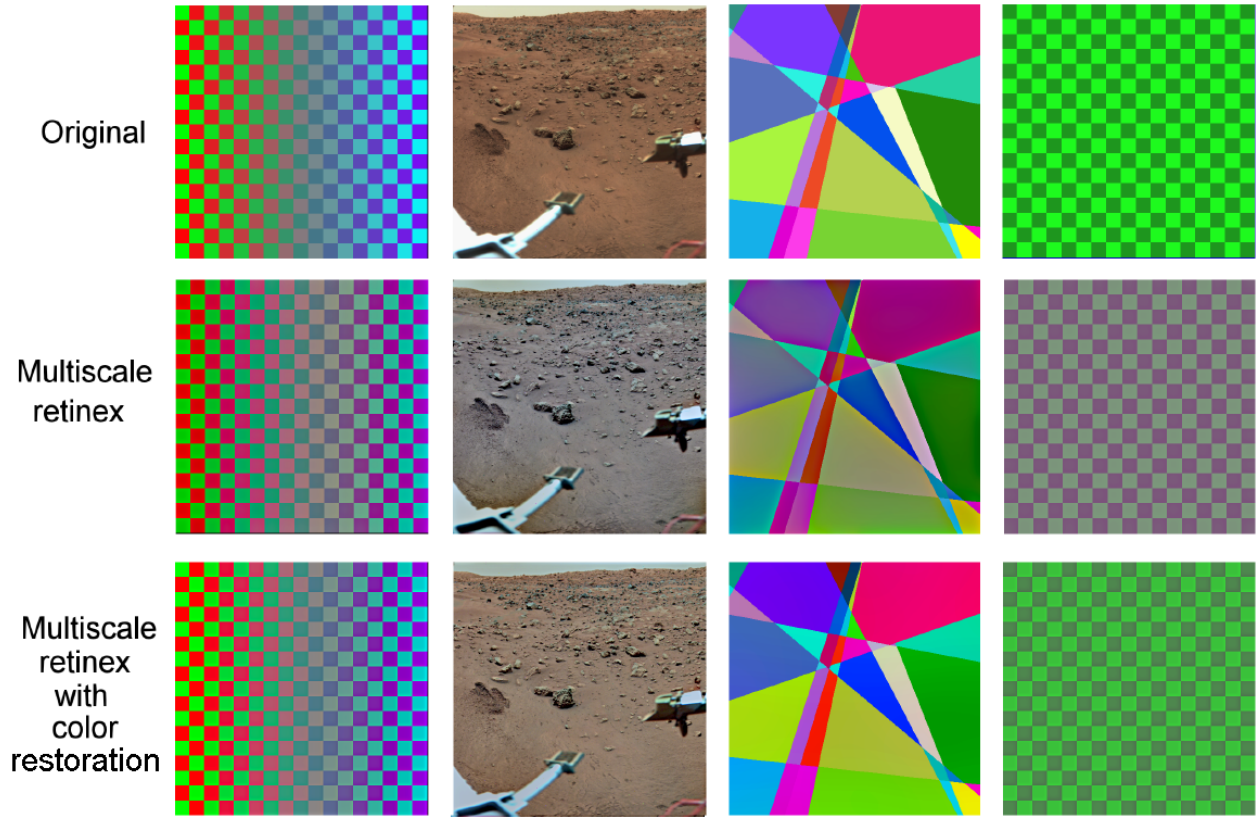

Figure 2. (Top row) Scenes that violate the gray-world assumption; (Middle row) the MSR output; note the graying of large areas of monochromes; (Bottom row) The MSRCR output; note that color constancy is diluted in order to achieve correct tonal rendition.

4. IMAGE ENHANCEMENT USING THE MSRCR

The MSRCR has been applied to literally thousands of images, both by us in a research environment, and by other users using the commercially available software package, PhotoFlair®* that implements the MSRCR. The MSRCR is capable of synthesizing color constancy, dynamic range compression, and contrast and lightness enhancement. Our experiments show that the MSRCR brings the *perception* of dark zones in recorded images up in local lightness and contrast to the degree needed to mimic direct scene viewing. Only images with very modest dynamic ranges do not need such enhancement and for these the exposure must be very accurate to achieve a good visual representation. Wide ranging reflectance values in a scene, and certainly, strong lighting variations demand a rather strong enhancement to achieve anything like the visual realism of direct observation. The dynamic range compression of the MSRCR is the basis for the contrast/lightness enhancement, and the generic character of the computation is the basis for using it as an automatic enhancement scheme. A few examples of retinex enhancements will serve to convey the degree to which images need to be improved and provide a demonstration that the MSRCR does, in fact, perform this task with considerable agility (Figures 3–5) and without human intervention. These examples highlight a major facet of retinex performance: intrinsically, the degree of automatic enhancement matches the degree of visual deficit in the original acquired image. Simply put, the MSRCR appears to mimic human perception in producing color and lightness that are influenced by the visual setting in which they occur. The exchange of spatial and spectral lighting dependencies for spatio-spectral context effects appears to be a very basic element of human vision and the MSRCR computation.

In addition to illustrating the agility of the MSRCR to automatically match the strength of the enhancement

*PhotoFlair®is a trademark of TruView Imaging Company `http://www.truview.com`

Figure 3. Retinex examples to illustrate that the strength of the enhancement matches the degree of visual deficit in the original image: original image (left column); subtle enhancements (right column).

to the degree of visual deficit, Figures 3–5 also illustrate the point made in Section 1 about the impact of the MSRCR processing on overall noise in the original image. In Figure 3 the MSRCR is applied to images that have good SNR across the whole image. The overall impact of the MSRCR is very subtle, increasing primarily the sharpness in certain areas of the image. The overall impression of the output image is a sharper image with slightly better contrast. Since the SNR is good across the original image, noise artifacts are not an issue in the MSRCR processed image.

Figure 4 shows the instance when the SNR is moderate, i.e., there are small areas in the image where there are shadows or lower intensity values. In such areas, the relative SNR is lower than in the rest of the image—signal strength is weaker but the noise strength is uniform across the image. Therefore, noise is enhanced and can be seen in the processed image in the areas where the original image contains shadows. The noise is on the verge

Figure 4. Retinex examples to illustrate that the strength of the enhancement matches the degree of visual deficit in the original image: original image (left column); moderate enhancements (right column).

of being objectionable but is acceptable given the amount of detail that is now visible. The overall impression is of a brighter, sharper image with more contrast.

In Figure 5, however, there are large areas that are dark and which possess poor contrast, and hence have poor SNR. The MSRCR processing enhances the noise in these areas since it treats the noise features as fine spatial detail. Depending upon the application, these enhancements may be acceptable or not. While the noise artifacts are strong in the processed image, so are the spatial details that were not visible in the original image. This is especially true in the top row of Figure 5 where the original image contains almost no details, while the MSRCR processed image is noisy but a significant amount of detail is present.

While the strength of detail enhancement produced by the MSRCR clearly matches the need for that enhance-

Figure 5. Retinex examples to illustrate that the strength of the enhancement matches the degree of visual deficit in the original image: original image (left column); strong enhancements (right column).

ment, it may not always be aesthetically pleasing to apply the MSRCR. We, however, are not really concerned about the aesthetic aspects of image enhancements—at least, not completely. In the next section we will examine the impact of noise on the performance of the MSRCR using a simulation environment that will allow us to control the amount of noise that is injected into the processing stream.

5. IMPACT OF RELATIVE NOISE ON THE PERFORMANCE OF THE MSRCR

Image quality is often quantified by the SNR or some variant of the SNR. The SNR is a global (scalar) measure where the whole image is characterized by a single number. However, the SNR across an image can vary dramatically especially in the presence of shadows where the SNR is considerably poorer than in areas of bright

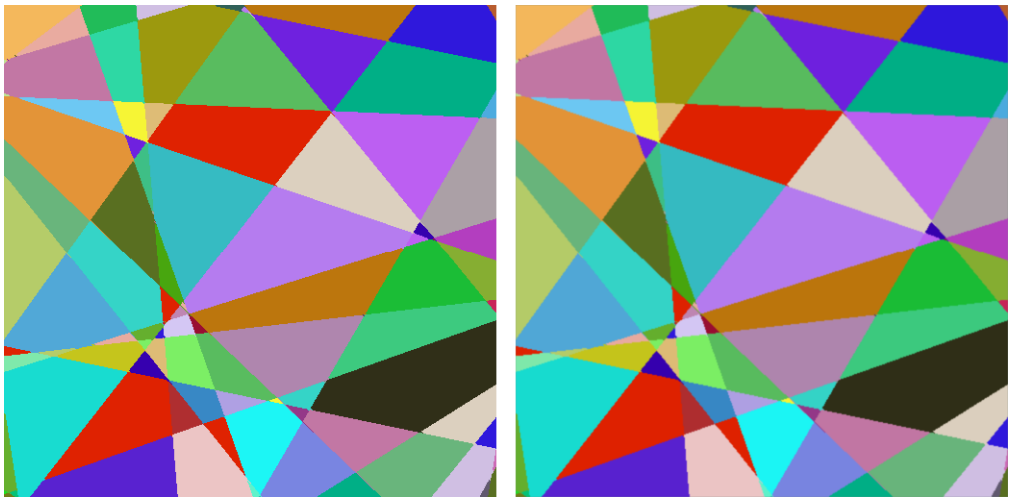

Figure 6. Original random polygon image (left); blurred version (right); the blurring simulates the impact of the lens on image acquisition.

illumination. This can be explained if one assumes that the noise that is injected into the image is white Gaussian. White Gaussian noise has a constant magnitude in the Fourier frequency. Hence it impacts different features in an image differently. The impact of this noise is strongest where the signal is the weakest, i.e., areas with poor illumination. This is exactly the scenario that we are interested in analyzing. For this purpose, we will use the simulation image shown in Figure 6 (left). The image in Figure 6 (right) is a slightly blurred version of the original image, where the blurring factor has been set to eliminate all aliasing.[†] We blur the original computer-generated image to mimic the process of acquiring an image of a scene with a digital camera. Since we start with a completely noise-free image, we can inject controlled noise into the process. Additionally we can project shadows on the simulated image. The combination of the two processes allows us to examine the impact of changing the shadow strength, hence the relative SNR, on the MSRCR process.

Figure 7 shows the shadow profiles for shadow strength $\mathcal{R}$ that depicts the strength of the shadow, $M_{\mathcal{R}}$, at the top-left and the bottom-left corners of the image. In other words, the signal strength weakens moving from the top-left of the image to the right and bottom, being weakest at bottom-right. Accordingly, since the strength of the noise is constant, the SNR is poorest at the bottom-right.

The noise free "shadowed" images and their MSRCR processed versions are shown in Figures 8. The MSRCR almost completely removes the impact of shadows $\mathcal{R} = 2, 4,$ and 8 and considerably mitigates the impact of the shadow for $\mathcal{R} = 16$. Figure 9 shows the original and MSRCR processed images when a white Gaussian noise that make the overall RMS SNR $= 8$ is added to the images. With the shadow, the relative SNR is considerably less than 8 where the shadow is the deepest.

In Figures 8 and 9, the impact of MSRCR processing is clearly evident to the eye. We want to quantify these results and for that we use the fidelity metric, $\mathcal{F}$. The fidelity between two $N_1 \times N_2$ images I_1 and I_2 is defined as:

$$\mathcal{F}(I_1, I_2) = 1 - \frac{\frac{1}{N_1 N_2} \sum_{x_1=0}^{N_1-1} \sum_{x_2=0}^{N_2-1} \left(I_1(x_1, x_2) - I_2(x_1, x_2)\right)^2}{\frac{1}{N_1 N_2} \sum_{x_1=0}^{N_1-1} \sum_{x_2=0}^{N_2-1} I_1(x_1, x_2)^2}, \tag{7}$$

and it measures how closely two images resemble each other. Figure 10 shows the fidelity $\mathcal{F}$ as a function of the shadow strength $\mathcal{R}$ for both the noise free imaging environment, and the imaging environment where the SNR $=$

[†]We will discount aliasing as a noise source in the rest of the paper.

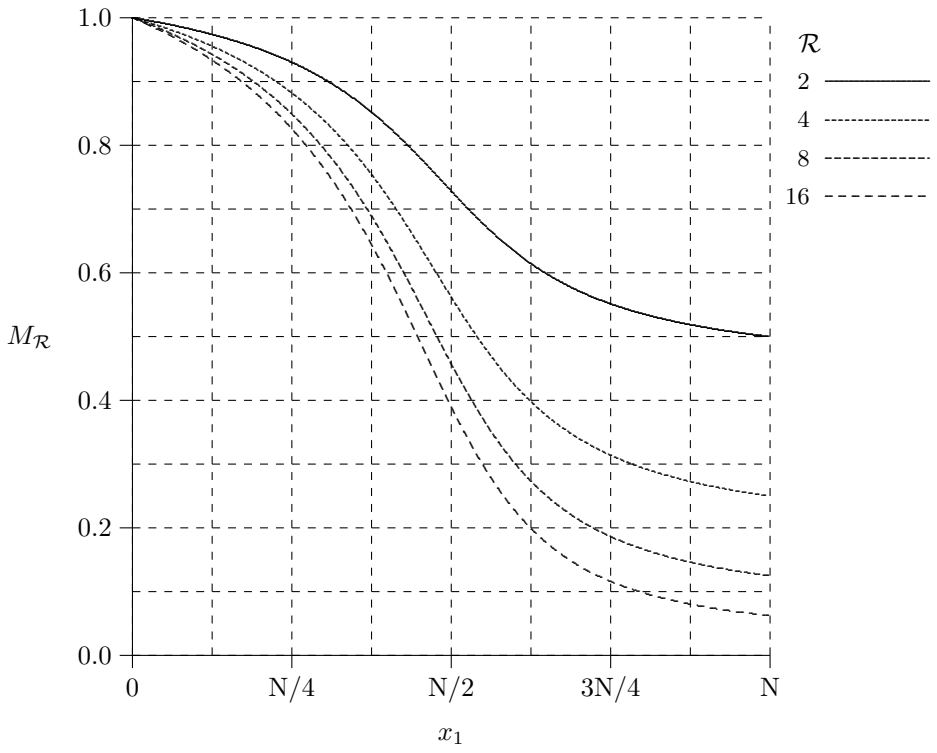

Figure 7. The shadow profile for $\mathcal{R} = 2, 8, 4$ and 16. $M_{\mathcal{R}}$ is the multiplicative factor at position x_1 (x_2) used to modify the image intensity $I_i(x_1, x_2)$.

8 for all images. The solid line in Figure 10 (top-left) shows the fidelity, $\mathcal{F}$, between the original image shown in Figure 6 (right) and the images shown in Figure 8 (top row), and the dashed line shows the same for the MSRCR output shown in Figure 8 (bottom-row). Similarly, the solid line in Figure 10 (top-right) shows $\mathcal{F}$ between the original image shown in Figure 6 (right) and the images shown in Figure 9 (top row) and the dashed line shows the same for the MSRCR output shown in Figure 9 (bottom-row). The bottom row in Figure 10 shows the relative improvement in $\mathcal{F}$ due to MSRCR processing as a function of the shadow strength. The relative improvement is given by $100\,(\Delta\mathcal{F}/\mathcal{F})$, where $\Delta\mathcal{F} = \mathcal{F}_m - \mathcal{F}_o$, $\mathcal{F}_o$ is the fidelity shown with the solid line in Figure 10, and $\mathcal{F}_m$ is the fidelity shown with the dashed line.

There are a few points that need to be made about the behaviors depicted in Figure 10:

1. The effectiveness of the MSRCR in eliminating shadows decreases as the shadow strength increases. Though this result is intuitively obvious, the reader should realize that this is more a short-coming of the imaging process than it is one of the MSRCR! The images used in this simulation are all generated with 8-bit quantization per color band. As the strength of the shadow approaches its minimum value—approximately 0.05 for $R = 16$—intensity values will range between 0–15. Hence there is very little signal strength in areas of deep shadow so the MSRCR is unable to recover it. If the quantization was set so that a wider dynamic range was available, then the performance of the MSRCR would improve accordingly.

2. Even with 8-bit quantization, the MSRCR processed images, in the absence of noise, have fidelity values ranging from 0.99 to 0.9 while the original images have fidelity values ranging from 0.9 to 0.55. This was clear to the eye in Figure 8 and the fidelity numbers back up the observation.

3. The addition of noise impacts the fidelity values for the MSRCR processed images more than it impacts the fidelity values for the original image. Figure 10 (top-right) shows this result. The fidelity values for MSRCR processed images now range between 0.94 and 0.8, while those for the original image vary between

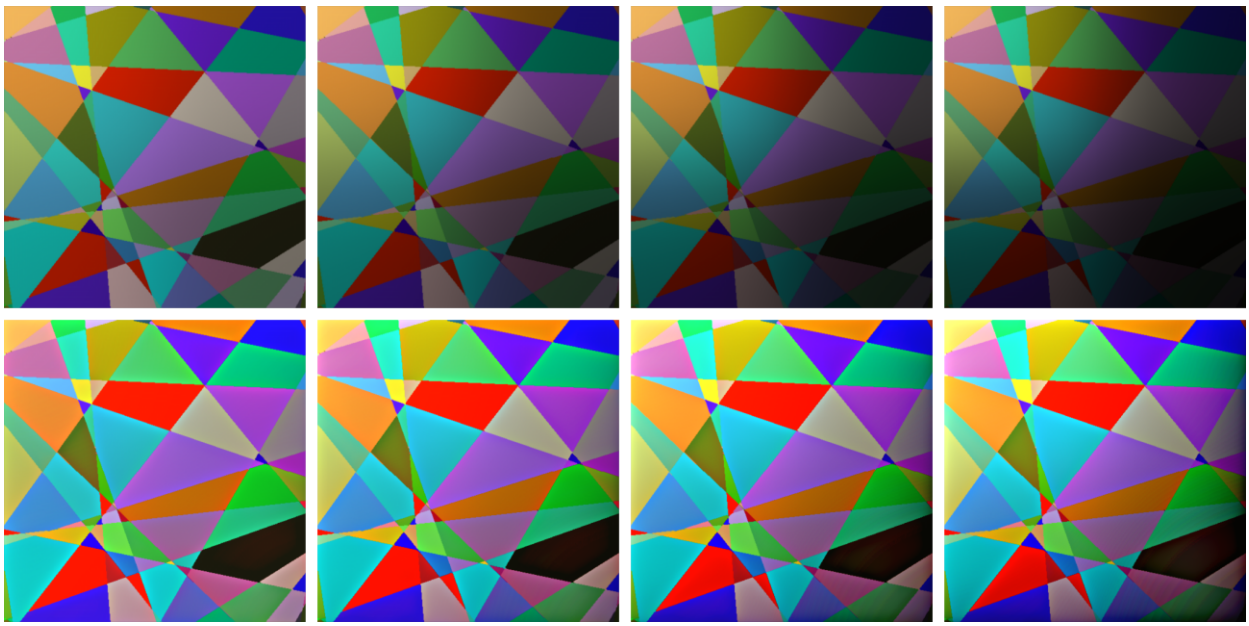

Figure 8. The image shown in Figure 6 (right) modified by the shadow profile shown in Figure 7 for $\mathcal{R} = 2, 4, 8, \text{and} 16$. The MSRCR processing either completely eliminates or considerably mitigate the impact of shadows.

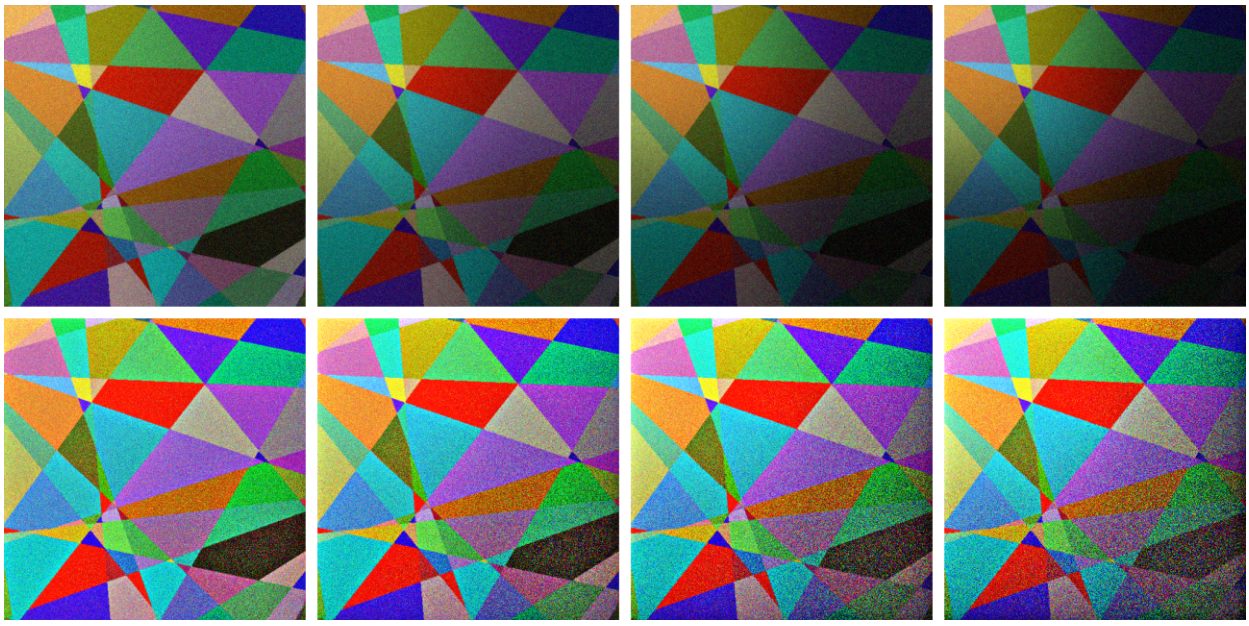

Figure 9. The original and MSRCR processed version of Figure 8 where white Gaussian noise has been injected into the images. The overall SNR for all the images is 8.

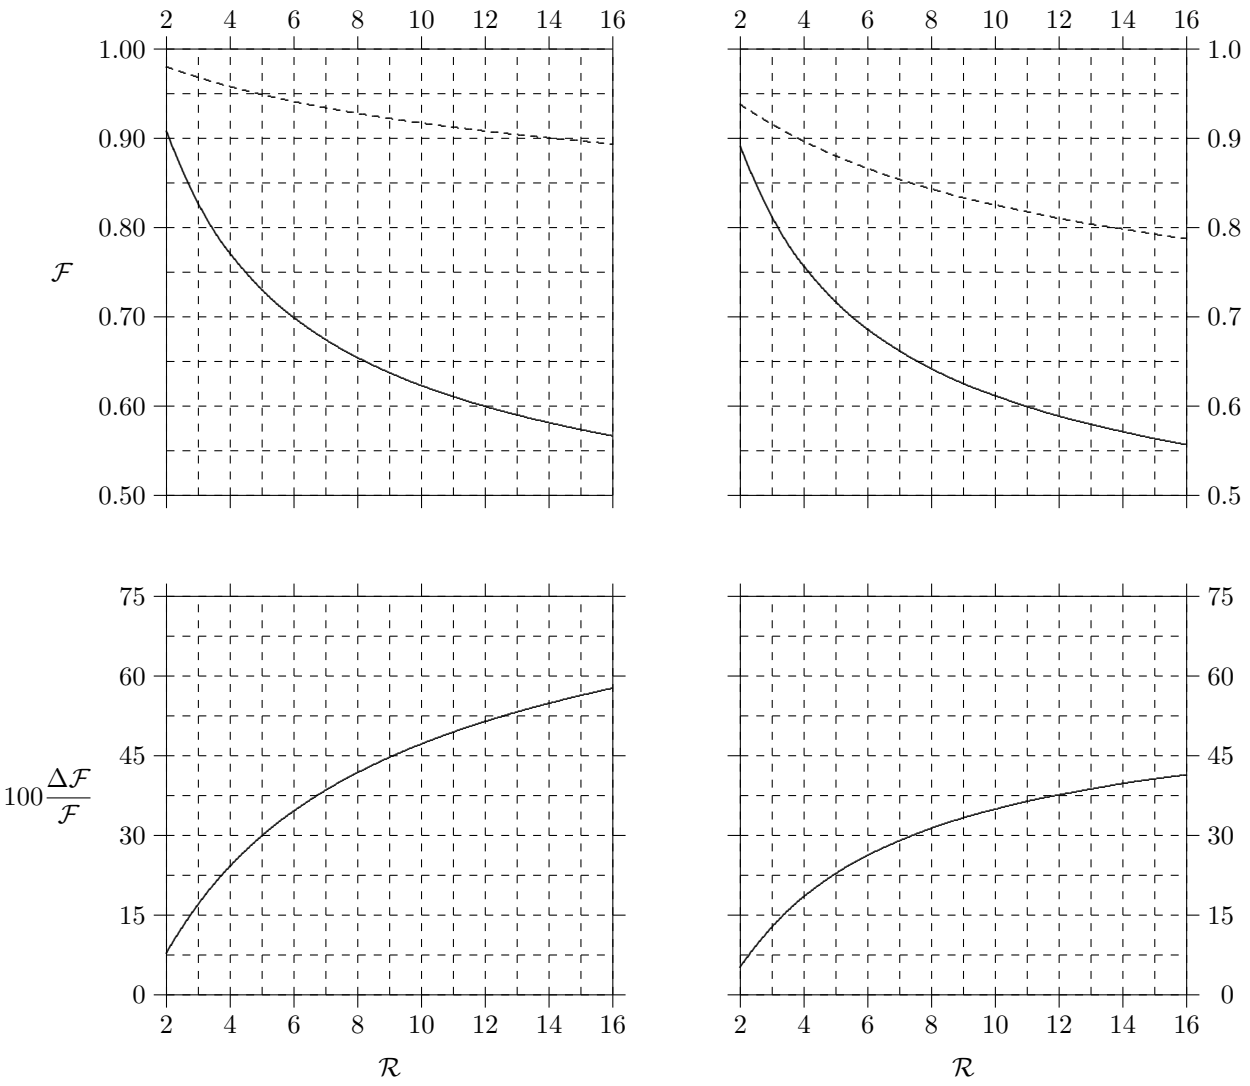

Figure 10. Fidelity, $\mathcal{F}$ as a function of the shadow strength $\mathcal{R}$. $\mathcal{F}$ for the image shown in Figure 6 (right). The left column shows $\mathcal{F}$ with respect to noise-free images shown in Figure 8. The right column shows $\mathcal{F}$ with respect to noise-added images shown in Figure 9.

0.55 and 0.90. So while there is a very slight change in $\mathcal{F}_o$, there is a significant change in $\mathcal{F}_m$. Figure 10 (bottom-row) also makes this point.

4. The relative increase in fidelity quantifies the assertion made in Section 4 that the the strength of enhancement produced by the MSRCR matches the visual deficit in the original image. As Figure 10 shows, as the $\mathcal{R}$ increases, so does $\Delta\mathcal{F}$. Since $\Delta\mathcal{F}$, in essence, measures the difference between the original data and the MSRCR processed data, we see that stronger deficits cause a stronger MSRCR output.

6. NOISE REDUCTION

In Section 5 we examined the impact of MSRCR processing on images that contain the same amount of noise but varying strength of shadows. The impact of shadows is to change the SNR across the image so different parts of the image have a different effective SNR. In this section we will examine the impact of varying the SNR on

the MSRCR when the original image does not contain shadows. We will use a darker image than the one shown in Figure 6 so that the impact of noise is clearly visible. We have previously examined aspects of this problem where we looked at the impact of MSRCR on segmentation of noisy images.[11]

Examination of the impact of noise on digital image data has been an active topic of research through the years. The techniques for noise reduction and noise amelioration have typically been based on mathematical analysis of the systems that are used to generate images, e.g., scanners and digital cameras. System analysis allows one to develop image *restoration* filters[12,13] that take into account all the different noise sources and minimize their combined impact with a Wiener-like filter. Other researchers[14–18] have used model based mathematical analysis to examine the impact of noise on the restorability of features but without relying heavily on the actual system design. These techniques are all very useful but have significant shortcomings in the presence of a high level of noise.

In previous papers[11,19,20] we have examined the issue of noise and its impact on the visibility of features in a different way. Instead of examining mathematical aspect of noise reduction, we made use of edge pattern analysis, both for automatic assessment of spatially variable noise and as a foundation for new noise reduction methods.[19] Starting with the premises that the edges in the image should be preserved, and that the overall impact of noise is to reduce resolution, we use connectivity analysis to find all those edge pixels in an image that are unconnected to any other pixel in a given neighborhood. We classify these edge pixels as "noise" and replace them by an average of their neighbors, hence reducing the resolution at the noise location. The impact of this operation is to leave the resolution of all connected edges intact while reducing resolution in areas where noise is detected. Since the overall impact of noise is to reduce resolution, reducing resolution in noisy areas by blurring eliminates the *appearance* of noise while not affecting overall image quality.

We have published this algorithm and some underlying issues that relate to this algorithm.[6] We however, repeat the procedure here for completeness:

1. Perform edge detection on the image using one of several edge-detection algorithms.
2. For each non-zero pixel in the edge image, perform connectivity analysis: check if the pixel is *connected* to another pixel in a 3×3 neighborhood; if it is, perform a similar analysis for that pixel, with the proviso that connectivity with the first pixel is ignored. Two pixels, $p(i,j)$ and $p(k,l)$ are said to be connected if

$$|p(i,j) - p(k,l)| < \mathcal{T} \tag{8}$$

   where $\mathcal{T}$ is some pre-defined threshold. if an edge pixel is connected to another edge pixel such that there is another edge pixel in between them, then it is considered to belong to a feature and not to noise.
3. When a noise pixel is found, replace it with an average of its neighbors:

$$p(i,j) = \big(p(i+1,j) + p(i,j+1) + p(i-1,j) + p(i,j-1)\big)/4. \tag{9}$$

   Otherwise do not modify the pixel.

We use the DOG (difference-of-Gaussians) filter followed by zero crossing detection to form the edge-only image. An important aspect of this algorithm is that while the analysis is performed on the edge image, the actions are taken on the original image. This procedure finds the noise pixels in the image and replaces them by a reduced resolution, blurred representation of its neighbors.

Figure 11 shows the fidelity $\mathcal{F}$ as a function of the SNR for both the original and the de-noised image. Results are also shown for the cases where the MSRCR is applied to the output of the denoising algorithm and where the denoising is performed after the MSRCR has been applied to the noisy image.

Figure 12 shows the images used for performing the fidelity analysis shown in Figure 11. Visual judgment corroborates the fidelity analysis. In each case, the MSRCR processing aggravates the noise artifacts considerably. Clearly, the results suggest that noise reduction should be performed before the application of the MSRCR. As Figure 11 (left) shows, the denoise-then-MSRCR process always provides better fidelity than the MSRCR-then-denoise process. This is corroborated by observing the data in Figure 12

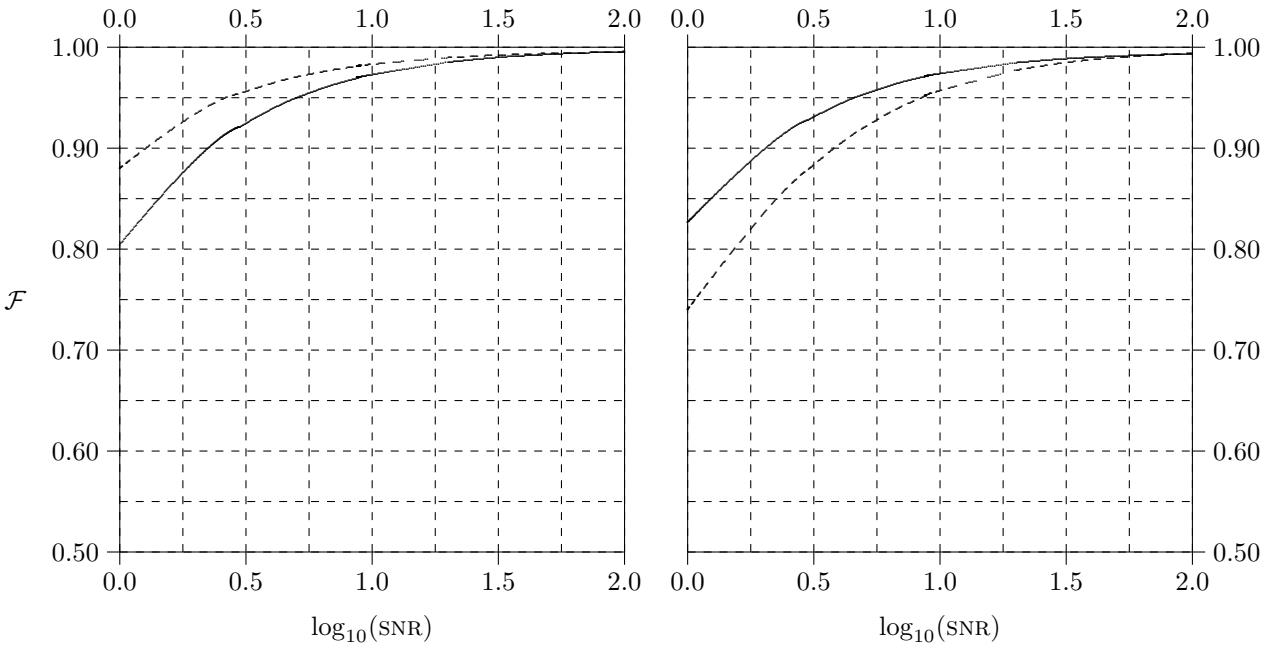

Figure 11. Fidelity, $\mathcal{F}$ as a function of the signal-to-noise ratio, SNR. (left) solid line shows the fidelity between original and noisy images, and the dashed line shows the fidelity between original and de-noised noisy images. (right) The solid line shows the fidelity between the original and the de-noised MSRCR processed image, and the dashed line shows the fidelity with the MSRCR processed, de-noised image.

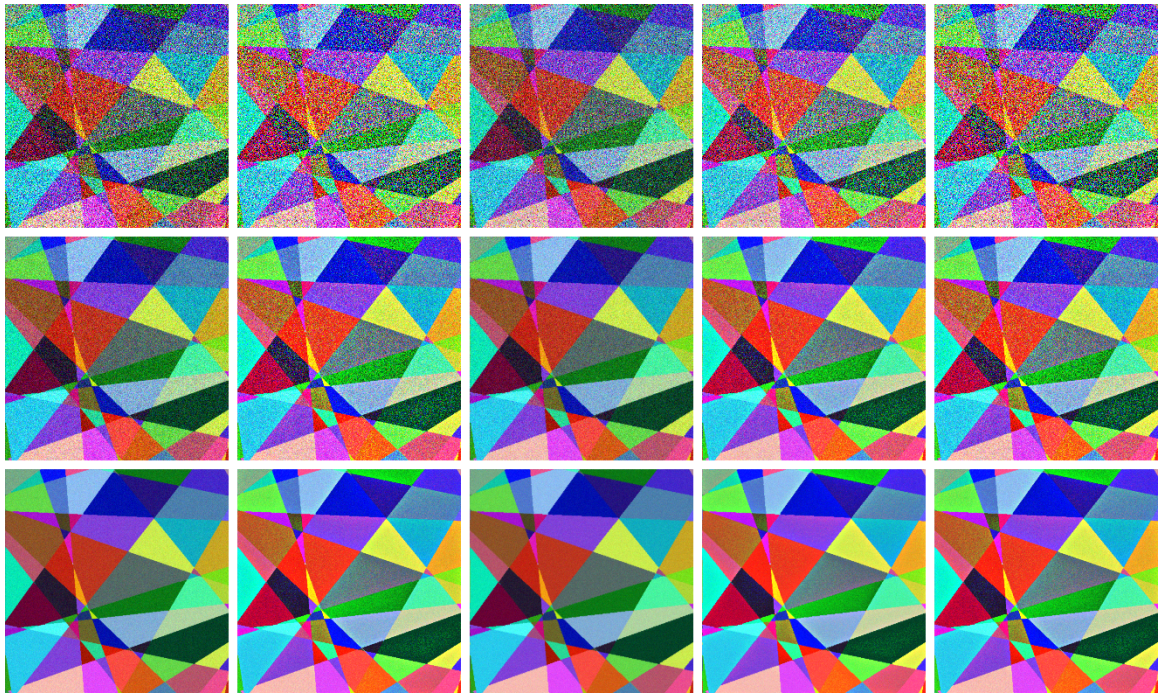

Figure 12. Noise and MSRCR processing. First column shows the noise added data for SNR = 1 (top), 10 (middle), and 100 (bottom). Second column shows the MSRCR outputs. Third column shows the denoised data where the original image has been de-noised based upon the algorithm outlined above. Fourth column shows the MSRCR applied to the de-noised data. And the last column shows the output of the denoising algorithm applied to the MSRCR output shown in the second column.

7. CONCLUSIONS

The MSRCR is a general purpose image enhancement algorithm that has been shown to perform very well on a variety of images including those containing deep shadows. When an image is acquired with a device with high quality optics and electronics, the MSRCR can retrieve almost all the fine details from shadows, accompanied with only a slight loss of color. However, since these systems tend to be very expensive, such a situation is not very general. Most images are acquired with consumer grade cameras where there is often a mismatch between the lens and the CCD, and the electronics are a source of considerable thermal noise. In this paper, we have presented a new analysis of MSRCR processing, where we have related its performance to SNR both in terms of image fidelity, and actual observation. We have used simulated digital data because this enables us to control the artifacts that are injected into the image due to the image acquisition process. Our findings show that whereas the MSRCR is very good at eliminating shadows under extremely severe conditions, its performance suffers when the digital data has poor SNR. The SNR may be poor due to quantization artifacts or the noise induced into the signal due to the sensor electronics. De-noising the data before MSRCR processing provides a significant gain in fidelity both when compared to denoising the data after MSRCR processing, and to not de-noising the data at all. The de-noising algorithm used here is based on retaining edges in an image and is in very early stages of development. Its performance can be enhanced considerably by improving the connectivity analysis that is used to determine whether data at a location is noise or signal. This is work under development and we hope to report on this in a future paper.

Acknowledgments

The authors wish to thank the Synthetic Vision Sensors element of the NASA Aviation Safety and Security Program for the funding which made this work possible. In particular, Dr. Rahman's work was supported under NASA cooperative agreement NNL04AA02A.

REFERENCES

1. D. J. Jobson, Z. Rahman, and G. A. Woodell, "Properties and performance of a center/surround retinex," *IEEE Trans. on Image Processing* **6**, pp. 451–462, March 1997.
2. D. J. Jobson, Z. Rahman, and G. A. Woodell, "A multi-scale Retinex for bridging the gap between color images and the human observation of scenes," *IEEE Transactions on Image Processing: Special Issue on Color Processing* **6**, pp. 965–976, July 1997.
3. E. Land, "Recent advances in retinex theory and some implications for cortical computations," *Proc. Nat. Acad. Sci.* **80**, pp. 5163–5169, 1983.
4. E. Land, "An alternative technique for the computation of the designator in the retinex theory of color vision," *Proc. Nat. Acad. Sci.* **83**, pp. 3078–3080, 1986.
5. E. Land, "Recent advances in retinex theory," *Vision Research* **26**(1), pp. 7–21, 1986.
6. Z. Rahman and D. J. Jobson, "Noise, edge extraction and visibility of features," in *Visual Information Processing XIV*, Z. Rahman, R. A. Schowengerdt, and S. E. Reichenbach, eds., Proc. SPIE 5817, 2005.
7. Z. Rahman, D. J. Jobson, and G. A. Woodell, "Multiscale retinex for color rendition and dynamic range compression," in *Applications of Digital Image Processing XIX*, A. G. Tescher, ed., Proc. SPIE 2847, 1996.
8. Z. Rahman, D. J. Jobson, and G. A. Woodell, "Retinex processing for automatic image enhancement," *Journal of Electronic Imaging* **13**(1), pp. 100–110, 2004.
9. W. D. Wright, *The Measurement of Colour*, Hilger and Watts, London, second ed., 1958.
10. P. Lennie and M. D. D'Zmura, "Mechanisms of color vision," in *CRC Critical Reviews of Neurobiology*, vol. 3, pp. 333–400, 1988.
11. Z. Rahman, D. J. Jobson, G. A. Woodell, and G. D. Hines, "Impact of multi-scale retinex computation on performance of segmentation algorithms," in *Visual Information Processing XIII*, Z. Rahman, R. A. Schowengerdt, and S. E. Reichenbach, eds., Proc. SPIE 5438, 2004.
12. C. L. Fales, F. O. Huck, and R. W. Samms, "Image system design for improved information capacity," *Applied Optics* **23**(6), pp. 872–888, 1984.

13. F. O. Huck, C. L. Fales, J. A. McCormick, and S. K. Park, "Image-gathering system design for information and fidelity," *Journal of the Optical Society of America A* **5**(3), pp. 285–299, 1988.
14. G. L. Anderson and A. N. Netravali, "Image restoration based on a subjective criterion," *IEEE Transactions on Systems, Man, and Cybernetics* **6**(12), pp. 845–853, 1976.
15. D. T. Kuan, A. A. Sawchuk, T. C. Strand, and P. Chavel, "Adaptive noise smoothing filter for images with signal-dependent noise," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **7**(2), pp. 165–177, 1985.
16. B. E. A. Saleh, "Trade off between resolution and noise in restoration by superposition of images," *Applied Optics* **13**(8), pp. 1833–1838, 1974.
17. J. W. Woods and V. K. Ingle, "Kalman filtering in two dimensions: Further results," *IEEE Transactions on Acoustics, Speech and Signal Processing* **29**(2), pp. 188–197, 1981.
18. D. T. Kuan, A. A. Sawchuk, T. C. Strand, and P. Chavel, "Nonstationary 2-D recursive filter for speckle reduction," in *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, IEEE, 1982.
19. D. J. Jobson, Z. Rahman, G. A. Woodell, and G. D. Hines, "The automatic assessment and reduction of noise using edge pattern analysis in nonlinear image enhancement," in *Visual Information Processing XIII*, Z. Rahman, R. A. Schowengerdt, and S. E. Reichenbach, eds., Proc. SPIE 5438, 2004.
20. D. J. Jobson, Z. Rahman, and G. A. Woodell, "Feature visibility limit in the enhancement of turbid images," in *Visual Information Processing XII*, Z. Rahman, R. A. Schowengerdt, and S. E. Reichenbach, eds., Proc. SPIE 5108, 2003.